Model Checking Software Programs with First Order Logic Specifications using AIG Solvers

Fadi A. Zaraket, Mohamad Noureddine

arXiv:1409.6825v1 [cs.SE] 24 Sep 2014

Abstract

Static verification techniques leverage Boolean formula satisfiability solvers such as SAT and SMT solvers that operate on conjunctive normal form and first order logic formulae, respectively, to validate programs. They force bounds on variable ranges and execution time and translate the program and its specifications into a Boolean formula. They are limited to programs of relatively low complexity for the following reasons. (1) A small increase in the bounds can cause a large increase in the size of the translated formula. (2) Boolean satisfiability solvers are restricted to using optimizations that apply at the level of the formula. Finally, (3) the Boolean formulae often need to be regenerated with higher bounds to ensure the correctness of the translation.

We present a method that uses sequential circuits, Boolean formulae with memory elements and hierarchical structure, and sequential circuit synthesis and verification frameworks to validate programs. (1) Sequential circuits are much more succinct than Boolean formulae with no memory elements and preserve the high-level structure of the program. (2) Encoding the problem as a sequential circuit enables the use of a number of powerful automated analysis techniques that have no counterparts for other Boolean formulae. Our method takes an imperative program with a first order logic specification consisting of a precondition and a postcondition pair, and a bound on the program variable ranges, and produces a sequential circuit with a designated output that is $true$ when the program violates the specification. Our method uses sequential circuit synthesis reduction techniques to reduce the generated circuit, and then uses sequential circuit verification techniques to check the satisfiability of the designated output. The results show that our method can validate designs that are not possible with other state of the art techniques, and with bounds that are an order of magnitude larger.

Index Terms

Software verification, static analysis, Boolean satisfiability solvers, Hoare triplet

1 INTRODUCTION

The work in [1] takes a declarative formula ϕ in first order logic (FOL) with transitive closure and a bound on the universe of discourse and translates it to a sequential circuit expressed in VHDL. A sequential circuit is a Boolean netlist with memory elements and a hierarchical structure. It then passes the sequential circuit to SixthSense [2], an IBM internal sequential circuit verification framework, and decides the validity of ϕ within the bound. It scales to bounds larger than what is possible with Kodkod [3] which translates ϕ into a propositional Boolean formula in conjunctive normal form (CNF) and checks its validity using a Boolean satisfiability solver such as MiniSat [4].

The work in [5] translates an imperative C program, with an assertion statement therein, and a bound on the input size, into a sequential circuit expressed in VHDL. It then passes the sequential circuit to SixthSense [2] and decides the validity of the assertion within the bound. It scales to bounds larger than what is possible with CBMC [6]. CBMC unwinds loops and recursive functions up to the given bound, and translates the program into a CNF formula that asserts the properties. It then checks for correctness using a Boolean satisfiability solver.

In this work, we present our method that takes an imperative program $\mathcal{S}$ with a specification, FOL precondition and postcondition pair (P, Q), and checks whether $\mathcal{S}$ satisfies the specification within a bound b on the domain of the program and specification variables ($\mathcal{S} \models (P, Q)|_b$); i.e. when the bounded inputs of $\mathcal{S}$ satisfy P, the outputs of $\mathcal{S}$ satisfy Q. The program is written in J, a subset of C++/Java that includes integers, arrays, loops, and recursion. Our method translates the problem $\mathcal{S} \models (P, Q)|_b$ into a sequential circuit with a designated output therein that is $true$ iff the program violates the specification within the bounded domain. Our method uses sequential circuit synthesis reduction and abstraction techniques embedded in the open source ABC [7] framework to reduce the generated sequential circuit. Then it uses ABC verification techniques to decide the satisfiability of the designated output. ABC either (1) proves the validity of the program, (2) generates a counterexample illustrating that the program violates

---

- *Fadi A. Zaraket is an Assistant Profesor with the Department of Electrical and Computer Engineering, American University of Beirut, Beirut, Lebanon.*
  E-mail: fadi.zaraket@aub.edu.lb
- *M. Noureddine is a graduate student at the Department of Electrical and Computer Engineering, American University of Beirut, Beirut, Lebanon.*
  E-mail: man17@aub.edu.lb

the specifications, or (3) reports an inconclusive result as it exhausts computational resources. Our method translates the counterexample back to the program domain and provides the user with a visual debugging tool (GTKWave [8]) to trace the violation.

Our method significantly extends the work of both [1], [5] in that

- it uses a program counter semantics to translate the program into an intermediary *one loop program* ($\mathcal{O}lp$), where a program counter is an additional variable that encodes the control flow of the program. It encodes the data flow of the program into ternary conditional statements based on the value of the program counter. The work in [5] performs only a source to source translation to VHDL,
- it directly translates the $\mathcal{O}lp$ program into bit level representation using *And-Inverter-Graphs* (AIG) sequential circuits while [1], [5] depend on the VHDL compiler and synthesis tools to do the translation to bit level,
- it supports imperative programs annotated with FOL specifications and also supports array boundary and overflow checks,
- it supports function calls including recursion,
- in case the original correctness check was not conclusive, it uses heuristics to guess a termination bound on the program execution time, it enables a termination guarantee check within the execution time bound, and it then uses the execution time bound with bounded model checking to decide correctness,
- it uses ABC [7], an open source sequential circuit synthesis and verification framework, instead of SixthSense [2] an IBM internal sequential circuit solver. ABC is a transformation-based verification (TBV) [7] framework that operates on sequential circuits and iteratively and synergistically calls numerous reduction and abstraction algorithms such as retiming [9], redundancy removal [10]–[13], logic rewriting [14], interpolation [15], and localization [16]. These algorithms simplify and decompose complex problems until they become tractable for ABC verification techniques such as symbolic model checking, bounded model checking, induction, interpolation, circuit SAT solving, and target enlargement [2], [17]–[20],
- our method is fully implemented as an open source tool ($\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$) available online [1].

We evaluated our method with the verification of standard algorithms, fundamental and complex data structures, real applications and programs from the software verification benchmarks

1. http://research-fadi.aub.edu.lb/dkwk/doku.php?id=sa

and compared our method to CBMC and other tools ranking top in the software verification competition [21]. $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ succeeded to find and report counterexamples for all defected programs. It found and reported defects that we were not aware of while developing the evaluation benchmarks. It scaled to verification bounds higher than those possible with the other tools, and proved specifications that were not possible by the other tools.

Limitations of translation to Boolean formulae: Recent advances in SAT enabled tools like Kodkod Alloy Analyzer [3], and CBMC [6] to check real programs. However, these programs often need to be partial, leaving out important functionality aspects, to enable the analysis to complete. Moreover, the analysis is typically bound to relatively small limits.

There are three limiting aspects of translating high-level programs to Boolean formulae.

Disadvantage 1 The translation to Boolean formulae depends on the bounds; a small increase in the bound on variable ranges can cause a large increase in the size of the translated formula due to unwinding loop and recursion structures in programs, or eliminating quantifiers in declarative first order logic.

Disadvantage 2 CNF SAT solvers are restricted to using optimizations, such as symmetry breaking [22] and observability don't cares (ODC) [23], that apply at the level of CNF formulae. However these optimizations usually aim at increasing the speed of the solver and often result in larger formulae as they add literals and clauses to the CNF formula to encode symmetry and ODC optimizations [24]. Often times when the analyzer successfully generates a large CNF formula, the underlying solver requires intractable resources.

Disadvantage 3 Often times the formula needs to be regenerated with higher bounds in case the unwinding bounds were not large enough for the loops to complete as is the case with CBMC and ESBMC.

To extend the applicability of static analysis to a wider class of programs as well as to check more sophisticated specifications and gain more confidence in the results, we need to scale the analysis to significantly larger bounds.

Advantages of sequential circuits: We formally define sequential circuits in Section 2.2; for now a sequential circuit can be viewed as a restricted C++ program, specifically a concurrent program in which all variables are either integers, whose range is statically bounded, or Boolean-valued, and dynamic allocation is forbidden [25]. There are two key advantages to compiling

programs into sequential circuits rather than Boolean formulae:

Advantage 1 Sequential encodings are much more succinct than SMT or pure combinational SAT formulae. They are imperative and state-holding while CNF and SMT formulae are declarative and state-free. For example, they can naturally represent the execution of quantifiers and loops without the need for unrolling them. Moreover, they can store and reuse intermediate results in local variables. In cases, SAT and SMT encoding algorithms produce a data structure that uses several orders of magnitude more memory to represent.

Advantage 2 Casting the decision problem for a program specification as an invariant check on a sequential circuit allows to leverage a number of powerful automated analysis techniques that we discuss in Section 2.3 and that have no counterpart in CNF or SMT analysis.

Other software verification techniques and tools exist that leverage predicate abstraction, interpolation, model checking, SMT, and other FOL and CNF solvers [26]–[32]. We discuss the tools and further compare our method to them in Section 7.

The rest of this paper is structured as follows. Section 2 reviews basic definitions and concepts related to Boolean formulae, and introduces the $Jcore$ and J programs, sequential circuits, and the ABC framework. Section 3 provides an overview of the $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$, introduces $\mathcal{O}lp$ programs, and illustrates the method with an array search example. Section 4 describes the translation of J programs into sequential circuits. Section 5 discusses the implementation. We discuss the results in Section 6, the related work in Section 7, and conclude with future work in Section 8.

2 Background

In this section we define programs, sequential circuits, and introduce the ABC synthesis and verification framework. A reader well-versed in software verification may wish to skip this section, using it only as a reference.

2.1 J programs

The grammar on the left of Figure 1 defines $Jcore$, the core subset of J. A $Jcore$ program is one or more declaration statements, followed by one or more statements. The directives `a-un-op`, `a-bin-op`, `b-un-op`, `b-bin-op`, and `ba-bin-op` denote arithmetic unary and binary operators, Boolean unary and binary operators, and Boolean arithmetic operators, respectively. The variables matching the `var` and `array-var` rules in a program $\mathcal{S}$ form the sets of scalar variables $V = \{v_1, v_2, \ldots, v_m\}$ and array variables $A = \{a_1, a_2, \ldots, a_n\}$, respectively.

```
program: declaration-statement+ statement+
statement: assignment | conditional | loop
list-of-statements: statement*

// statements
assignment: target = expression ;
conditional: if (boolean-expr) { then-block } else { else-block }
loop: while ( boolean-expr ) { while-block }
then-block, else-block, while-block: list-of-statements

// declarations
declaration-statement: declaration ;
declaration: (variable-decl | array-decl)
variable-decl: type var
array-decl: type array-var [ constant? ] ([ constant ])?
var, array-var : id
type: int | bool

// expressions
boolean-expr: term | b-un-op boolean-expr | boolean-expr b-bin-op
     boolean-expr | expression ba-bin-op expression
term: constant | target-term
target-term: var | array-access
array-access: array-var [ expression ] | array-var [ expression ]
    [ expression ]
expression: term | a-un-op expression | expression a-bin-op
    expression | boolean-expr ? expression : expression
```

```
// statement extended with pre/post conditions
statement: assignment | conditional | loop | pre-condition
    | post-condition
// specification pre and postcondition
pre-condition: @pre specname { boolean-expr }
post-condition: @post specname { boolean-expr }

// boolean expression extended with quantifiers
boolean-expr: term | b-un-op boolean-expr | boolean-expr
    b-bin-op boolean-expr | expression ba-bin-op expression
     | quantifier
quantifier: (forall|exists) ( range ) { boolean-expr }
range: var [ expression .. expression ]

// declaration statement extended with function declaration
declaration-statement: declaration ; | function-decl
function-decl: type fname ( arg-decl-list? ) {
    declaration-statement* body-block }
arg-decl-list: variable-decl ( , variable-decl)*
body-block: list-of-statements return-statement
return-statement: return expression;

// terms extended with function calls
term: constant | target-term | function-call
function-call: fname ( arg-call-list? )
arg-call-list: expression ( , expression )*
specname, fname : id
```

Fig. 1: Grammar of *Jcore* (left), the core subset of the J imperative language (right).

Definition 1 (terms): A *target term* is either a variable $v \in V$, or an array access term of the form $a[e]$ which denotes the e^{th} element of a where $a \in A$ and e is an expression. A *term* is either a target term, or a constant $c \in \mathbb{Z}$.

Definition 2 (expressions): An *expression* is a term, an arithmetic expression of the form $-e, e_1 + e_2, e_1 - e_2, e_1 * e_2, e_1/e_2, e_1\%e_2$ where e, e_1, e_2 are expressions and $-, +, *, /$, and $\%$ denote subtraction, addition, multiplication, division and remainder, respectively.

Definition 3 (Boolean expressions): A *Boolean expression* is either (1) a constant from the set $\mathbb{B} = \{true, false\}$, (2) a binary Boolean arithmetic expression of the form $e_1 < e_2, e_1 \leqslant e2, e_1 > e_2, e_1 \geqslant e_2, e_1 == e_2$ where e_1, e_2 are expressions and $<, \leqslant, >, \geqslant$, and $==$ denote smaller, less than or equal, bigger than, bigger than or equal, and equal, respectively, (3) a binary Boolean expression of the form $b \&\& b', b \| b', b \rightarrow b', b == b'$, or (4) a unary Boolean expression of the form $\neg b$ where b, b' are Boolean expressions and $\&\&, \|, !, \rightarrow$, and $==$ denote logical conjunction, disjunction, negation, implication, and equivalence, respectively.

Definition 4 (First order logic formula): A *first order formula* is either a Boolean expression, or

a quantified formula of the form $Qq.b(q)$, where $Q \in \{\forall, \exists\}$ is a universal or an existential quantifier, q is a quantified variable, and $b(q)$ is a first order formula with q as a free variable.

Definition 5 (Statement): A *statement* is (1) an *assignment* statement of the form $t = e$ where t is a target term and e is an expression, (2) a *list of statements* in a well defined order that could be empty, (3) a *conditional* statement $if(b_1)\{\ then - block\}\ else\ \{\ else - block\}$ where b_1 is a Boolean expression, and $then - block$, and $else - block$ are lists of statements, (4) a while loop statement $while(b_2)\{while - block\}$ where b_2 is a Boolean expression and $while - block$ is a list of statements, or (5) a declaration statement of the form $type\ v;$ or $type\ a[constant];$ where $type$ is either `bool` or `int` denoting the domain of the variable values, v and a denote the variables in sets V and A, and `constant` denotes the size of the corresponding array.

The grammar of Figure 1 (right) extends $Jcore$ to form J. Boolean expressions are extended with existential and universal quantifiers that quantify a Boolean expression over a range of variable values defined by a start and an end expression. Statements are extended with pre- and postconditions that share the same specification name (`specname`). Declaration statements are extended with function declarations that take an argument declaration list, and a body block including a return statement. Terms are extended with function calls with an argument call list.

J program semantics: The semantics of a program is defined in terms of traces of variable values across execution time. The assignment statements define the data flow of the program and their semantics is defined in terms of updating the value of the target term with the value of the expression of the assignment statement. The lists of statements, conditional statements, loop statements, function calls, and return statements define the control flow of the program. The program starts execution at the first statement in its list of statements. The list of statements defines the order of execution. An if/else conditional statement s in a list of statements ℓ executes the `then-block` if the value of the Boolean expression b is $true$, otherwise it executes the `else-block`. The last statement in the `then-block` and the `else-block` moves execution to the statement next to s in ℓ. A loop statement s in a list of statements ℓ executes the first statement of the `while-block` if the value of the Boolean expression is $true$, otherwise it executes the statement next to s in ℓ. The last statement in the `while-block` moves execution back to the loop statement s.

A function call updates the values of the function argument declaration variables with the corresponding function call argument expressions and moves execution to the first statement of

the `body-block`. The return statement in a function declaration substitutes the corresponding function call with the value of the return expression and moves execution back to the statement that made the function call.

A specification (P, Q) is a pair of first order formulae where P is the precondition specifying constraints on the inputs of $\mathcal{S}$ and Q is a postcondition relating the outputs of $\mathcal{S}$ to its inputs.

2.2 AIG Sequential Circuit

Definition 6 (Sequential circuit): A *sequential circuit* is a tuple $\big((U, E), G, O\big)$. The pair (U, E) is a directed graph on vertices U and edges $E \subseteq U \times U$ where E is a totally ordered relation. The function $G : U \mapsto types$ maps vertices to $types$. There are three disjoint types: *primary inputs*, *bit-registers* (which we often simply refer to as *registers*), and logical *gates*. Registers have designated *initial value*, as well as *next-state functions*. Gates describe logical functions such as the conjunction or disjunction of other vertices. A subset O of U specifies the *primary outputs*. We denote the set of primary input variables by I, and the set of bit-register variables by R.

Definition 7 (Fanins): We define the direct *fanin*s of a gate u to be $\{v \mid (v, u) \in E\}$ the set of source vertices connected to u in E. We call the *support* of u $\{v \mid (v \in I \vee v \in R) \wedge (v, u) \in *E\}$ all source vertices in R or I that are connected to u with $*E$, the transitive closure of E.

For the sequential circuit to be syntactically well-formed, vertices in I should have no fanins, vertices in R should have 2 fanins (the next-state function and the initial-value function of that register), and every cycle in the sequential circuit should contain at least one vertex from R. The initial-value functions of R shall have no registers in their support. All sequential circuits we consider will be well-formed.

The ABC analyzer reasons about AIG sequential circuits which are sequential circuits with only NAND gates restricted to have 2 fanins. Since NAND is functionally complete, this is not a limitation.

Semantics of sequential circuits:

Definition 8 (State): A *state* is a Boolean valuation to vertices in R.

Definition 9 (Trace): A *trace* is a mapping $t : U \times \mathbb{N} \to \mathbb{B}$ that assigns a valuation to all vertices in U across time *steps* denoted as indexes from $\mathbb{N}$. The mapping must be consistent with E and G as follows. Term u_j denotes the source vertex of the j-th incoming edge to v, implying that

TABLE 1: Brief description of selected ABC synthesis and verification techniques.

Technique	Description	Command
Balancing [7]	Logic balancing applies associativity transformation to reduce AIG levels.	`balance`
Sweep	Structural register sweep (SRS) reduces the number of registers in the circuit by eliminating stuck-at-constant registers [34].	`ssweep`
Correspondence	Signal correspondence (Scorr) computes a set of classes of sequentially-equivalent nodes using k-step induction [34].	`scl -l`
Rewriting	AIG rewriting iteratively selects and replaces rooted subgraphs with smaller pre-computed subgraphs in order to reduce AIG size [35].	`rewrite`
Refactoring	Refactoring is a variation of rewriting. It uses a heuristic to compute a large cut for selected AIG nodes, then replaces the sub-graph that corresponds to the cut with a refactored structure if an improvement is observed [36].	`refactor`
Retiming	Retiming manipulates register boundaries and count in a given logic network, while maintaining output functionality and logic structure [37].	`retime`
Induction	Temporal induction uses circuit SAT and BDD solvers to carry simple and k-step induction proofs over the time steps of the AIG [38].	`ind`
Interpolation	Interpolation-based algorithms aim find interpolants and overapproximate the reachable states of the AIG with respect to the property [39].	`int`
Reachability	Property directed reachability (Pdr) tries to prove that there is no transition from an initial state of the AIG to a bad state [40].	`pdr`

$(u_j, v) \in E$. The value of gate v at time i in trace t is denoted by $t(v, i)$.

$$t(v,i) = \begin{cases} s_v^i & : v \in I \text{ with sampled value } s_v^i \\ t(u_1, 0) & : v \in R, i = 0, u_1 := \text{ initial-state of } v \\ t(u_2, i-1) & : v \in R, i > 0, u_2 := \text{ next-state of } v \\ G_v\big(t(u_1,i), ..., t(u_n,i)\big) & : v \text{ is a combinational gate with function } G_v \end{cases}$$

The semantics of a sequential circuit are defined with respect to semantic traces. Given an input valuation sequence and an initial state, the resulting trace is a sequence of Boolean valuations to all vertices in U which is consistent with the Boolean functions of the gates. We will refer to the transition from one valuation to the next as a *step*. A node in the circuit is satisfiable if there is an input sequence which when applied to an initial state will result in that node taking value $true$. A node in the circuit is valid if its negation is not satisfiable. We will refer to targets and invariants in the circuit; these are vertices whose satisfiability and validity is of interest, respectively. A sequential circuit can naturally be associated with a finite state machine (FSM), which is a graph on the states. However, the circuit is different from its FSM; among other differences, it is exponentially more succinct in almost all cases of interest [33].

2.3 ABC synthesis and verification framework

ABC is an open source synthesis and verification framework for sequential circuits. ABC operates on sequential circuits in AIG format and checks the satisfiability of a designated output gate

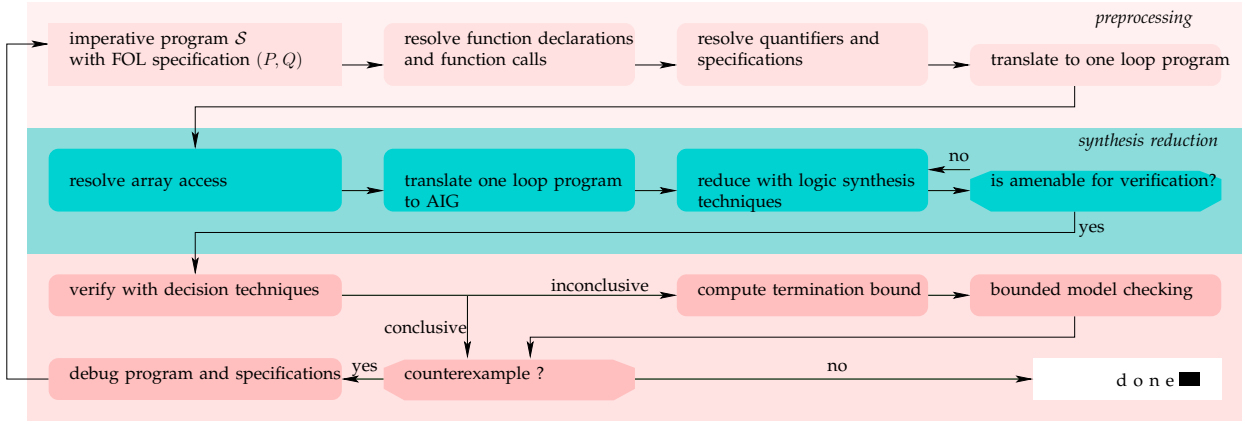

Fig. 2: Overview of verifying J programs with AIG synthesis and verification techniques

therein. ABC applies several reduction and abstraction techniques to simplify and decompose the problem into smaller problems. It then calls decision techniques to decide the simplified problems. Table 1 briefly summarizes some of the techniques supported by ABC.

3 Overview

Figure 2 illustrates our method. First, $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ preprocesses program $\mathcal{S}$ with its FOL specification pair (P, Q) to resolve function calls, specifications, and quantifiers and produces an equivalent program $\mathcal{S}'$ in $Jcore$. It then translates $\mathcal{S}'$ into a one loop program where it encodes the control flow into an additional program counter variable and encodes the data flow into assignment statements with ternary conditional expressions that depend on the value of the program counter variable. Then $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ translates the generated one loop program into an AIG circuit with bit vectors of fixed width that correspond to program variables. Array sizes are limited by the largest possible index and $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ translates each array element into a vector of AIG registers. $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ resolves array access operations to refer to the vectors of registers, it then instantiates equivalent logic circuits to the expressions and connects the circuits to the initial value and next state value functions of the registers. $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ designates one output of the AIG circuit to be $true$ when the program violates the specifications. Other AIG outputs signal out of bound array access and arithmetic overflow.

$\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ uses the ABC synthesis techniques to reduce the size and the complexity of the generated AIG circuit until it is amenable for verification. Once the AIG circuit is small enough, $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ uses ABC verification algorithms to check the designated output. The verification technique may find a violation and return a counterexample, return a proof that the program is

correct within the bound, or return an inconclusive result after exhausting computational limits. When an ABC verification returns a counterexample, $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ translates the counterexample from the AIG circuit level to traces of variable values in the original program and provides the user with a visual debugging view using GTKWave [8].

We describe the termination guarantee bound technique, the $\mathcal{O}lp$ programs, relate them to AIG circuits, and then illustrate the translation to $\mathcal{O}lp$ using an array search example.

3.1 Termination guarantee bound

In case the ABC prove techniques do not reach a conclusive result, $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ uses heuristics to compute a termination upper bound θ in terms of execution time, and calls ABC twice. The first call is to verify that the program is guaranteed to terminate within the termination bound ($\psi \equiv time >= \theta \rightarrow pc = last(\mathcal{S})$). The second call uses θ as a bound with the ABC bounded model checking technique to prove that the program does not violate the specification within the termination bound. Intuitively, checking $\mathcal{S} \models (P, \Psi)|_b$ is often easier than checking $\mathcal{S} \models (P, Q)|_b$, and once θ is confirmed as an upper bound for termination, bounded model checking is decidable.

3.2 One loop programs

The grammar in Figure 3 extends *Jcore* with the *primary* and *dotogether* constructs to define one loop programs. The `primary` construct denotes a primary input variable declaration with non-deterministic values. The concurrent *dotogether*{*together-block*} denotes that all statements of the *together-block* list of statements execute simultaneously. An $\mathcal{O}lp$ starts with variable declarations. Then the `init-list` list of assignment statements concurrently initializes the $\mathcal{O}lp$ variables. After initialization, a loop keeps updating the values of the variables concurrently using the list of assignment statements `next-list` until the $\mathcal{O}lp$ is done as denoted by the *notdone* Boolean variable. Intuitively, the structure of an $\mathcal{O}lp$ is similar to that of a sequential circuit where the variables

```
primary-input: primary declaration;

decl-list: primary-input*
    declaration-statement+
init-list: assignment+
next-list: assignment+

one-loop-program: decl-list
    @dotogether {
      init-list }

    while (notdone) { @dotogether {
      next-list } }
```

Fig. 3: $\mathcal{O}lp$ Grammar

```
1  int ArraySearch(int[] a,int d,int s,int e,int n) {
2   @pre as { 0<=s && s<=e && e<n && n<=MAXSIZE; }
3   int i = s;          // pc = pc+1
4   while ( i <= e) {   // pc = (i <= e) ? 5 : 10;
5    if (a[i] == d) {   // pc = (a[i] == d) ? 6 : 8;
6     break;}           // pc = 10;
7    else {
8     i = i+1;}         // pc = 4;
9    }
10  return i;           // pc = 11; rv = i;
11 @post as {
12  (s <= rv && rv <= e) -> a[rv] == d
13  forall(int i)[s .. e]{
14    a[i] != d -> (rv== -1) }
15 } }
```

```
1  @dotogether { // init-list for initial value0s
2   preas = 0 <= s && s <= e && e < n && n <= MAXSIZE;
3   pc = 0; notdone = true; postas = true; }
4  while (notdone) { @dotogether {
5   // next-list with next state functions
6   i = (pc == 3) ? s : (pc == 8) ? i+1 : i;
7   notdone = (pc == 11) ? false : true;
8   rv = (pc == 10) ? i : rv;
9   pc = (pc == 0) ? 3: (pc == 3) ? 4 :
10        (pc == 4) ? ( (i <= e) ? 5 : 10 ) :
11        (pc == 5) ? (a[i] == d) ? 6 : 8 ) :
12        (pc == 6) ? 10 : (pc == 8) ? 4  :
13        (pc == 10) ? 11 : pc;
14  postas = (rv >= s && rv <= e) -> a[rv] == d &&
15    forall(int i)[s .. e]{a[i] != d -> (rv== -1); }}}
```

Fig. 4: Array search in J and its equivalent $\mathcal{O}lp$ with the program counter variable.

correspond to registers, the `init-list` and `next-list` correspond to the initial and next state value functions, respectively.

3.3 Array search example

The Array search program in Figure 4 takes as input an array a, a start index s, an end index e, a data value d, and the number of elements in the array n. The precondition states that s and e are within array bounds and that n is within the bound of array sizes. The postcondition makes use of a variable rv denoting the return value. It states that if rv is valid between s and e inclusive, then $a[rv]$ must be equal to d, otherwise, rv must be invalid (-1) and all entries in a between s and e inclusive are not equal to d. Figure 4 also shows a semantically equivalent $\mathcal{O}lp$ array search program in terms of the values of the common program variables at termination.

The equivalent program introduces Boolean variables `preas`, `postas`, and `notdone` to encode the precondition, postcondition, and termination state of the program, respectively. The equivalent program also introduces a program counter variable `pc` which encodes the control flow of the program as indicated in the comments of the original program.Variable `notdone` is initialized to $true$, and `pc` is initialized to the first executable line of the program 3. Once `pc` reaches the last executable line of the program 13, the program terminates and thus `notdone` becomes $false$. Assignment statements are grouped by target variables, and encoded into ternary conditional expressions that depend on the value of `pc`. For example, the iterator i is assigned to s when `pc` is 3, incremented when `pc` is 8, and remains the same otherwise. Furthermore, the assignment statements on Lines 2 and 3 assign initial values to the target variables. The

assignment statements inside the `while` loop (Lines 5 to 15) compute the next state value of each of the target variables.

Our method translates the $\mathcal{O}lp$ program to an AIG sequential circuit where an iteration of the `while` loop is equivalent to a single time step in the AIG. For example, The `pc` variable ranges from 0 to 11 and is encoded with bit-vector of 4 registers with initial value functions set to 0 and next state functions set to gates representing the right hand side expression of the pc assignment statement. Section 4.4 discusses the translation from $\mathcal{O}lp$ to AIG.

Our method takes the resulting AIG and designates a gate therein that represents $\neg(preas \wedge done \rightarrow postas)$ as the output gate, passes the AIG to ABC, and checks for validity. The ABC solver returns the counterexample of Figure 13 with $a = [15\ 15\ 15\ 11], s = 3, e = 3, n = 4, d = 13, rv = 0$ where d is not in a, and the return value is $e + 1$, while the postcondition requires an invalid index (-1). The provided counter example can be used to fix the program. A possible fix is to replace Line 6 with `return i;`, and Line 10 with `return -1;`. Our method takes the fixed program, translates it into a sequential circuit, and passes it to ABC which validates the correctness of the fixed program using symbolic model checking.

4 TRANSFORMATION

In this section we present source to source transformation algorithms to translate a program $\mathcal{S}$ written in $Jcore$ into a $\mathcal{O}lp$ program. We then present algorithms that preprocess a J program with functions and specifications and translates it into $Jcore$. Finally, we present algorithms for translating the generated $\mathcal{O}lp$ program into an AIG circuit. The algorithms follow the execution semantics of J and the correctness of the translation holds by construction. We first present helper functions and terms that are used in the algorithms.

4.1 Helper functions

Function `label(`s`)` takes a statement in $\mathcal{S}$ and returns a unique label in $\mathbb{N}$. Function `next(`s`)` takes a statement s and returns the label of the next statement of s in a list of statements ℓ. If s is the last statement in ℓ, several cases arise.

1) ℓ is a function declaration `body-block`, `next` returns the label of the return statement.
2) ℓ is a `while-block` of a loop l, `next` returns the label of the loop statement `label(`l`)`.
3) ℓ is a `then-block` or an `else-block` of a conditional s_1, `next` returns the label of the statement next to s_1 (`next(`s_1`)`).

4) Finally, ℓ is the last statement in $\mathcal{S}$, `next` returns a special label $done$.

Functions `first` and `last` return the first and the last statement in a list of statements, respectively, or nil if the list is empty. Functions `then`, and `else` take a conditional statement s of the form `if (`b`) {then-block} else {else-block}`. They return the labels of the first statement of the then and else blocks, `label( first( then-block))`, and `label( first( else-block))`, respectively. If the list of statements is empty, both functions return `next(`s`)`.

Function `body` takes a loop statement or a function declaration. The loop statement is of the form `while (`b`) {while-block}`. The function declaration is of the form `type fname( arg-decl-list) { body-block }` where `type`, `fname`, and `arg-decl-list` are the return type, name, and argument declaration list of the function, respectively, and `body-block` is a list of statements followed by a return statement. For a function declaration, `body` returns the label of the first statement in `body-block` (`label(first(body-block))`). For a loop, `body` returns the label of the first statement in `while-block`. In case `while-block` is empty, it returns the label of the loop. Function `condition` takes a conditional or a loop statement and returns b.

Functions `target` and `expr` take an assignment statement of the form `target-term = expression;` and return `target-term` and `expression`, respectively. Function `base` and `index` take an array access term of the form `a[`e_1`]` or `a[`e_1`][`e_2`]` where a is an array variable, and e_1 and e_2 are expressions. Function `base` returns a. When a is declared as `type a[`c_1`]`, where c_1 is a constant, `index` returns e_1. When a is declared as `type a[`c_1`][`c_2`]`, where c_2 is a constant, `index` returns $e_1c_2 + e_2$. Finally, function `typeof` takes a variable (or array variable) and returns its type.

4.2 *Jcore* programs to $\mathcal{O}lp$

Algorithm `generateOneLoopProgram` of Figure 5 takes a program $\mathcal{S}$ written in *Jcore* and generates the `decl-list`, `init-list`, and `next-list` of an equivalent $\mathcal{O}lp$ program. Algorithm `generateDeclarationList` generates a declaration for each variable in $\mathcal{S}$, and also generates a primary input `vnondet` variable for each variable that is used before being initialized in $\mathcal{S}$. It also declares the program counter pc and the termination guard $notdone$ variables. Algorithm `generateOneLoopProgram` calls `generateInitList` that generates `init-list` where all variables are initialized to 0 except those that are used before being assigned in $\mathcal{S}$, those are initialized with their corresponding non-deterministic primary input variables. The program counter pc and the termination guard $notdone$ are initialized to the label of the first statement of the program, and to $true$, respectively.

```
generateOneLoopProgram(S)
 generateDeclarationList(S)
 generareInitList(S)
 // translate data flow into next-list
 foreach variable v in V ∪ A
  if v ∈ V
   v−next−list = generateVarNextList(v,S)
  else // v ∈ A
   v−next−list = generateArrayNextList(v,S)
  endif
  append v−next−list to next−list
 endfor
 // translate control flow into next-list
 pc−next−list = generatePCNextList(S)
 append pc−next−list to next−list
 append notdone = !(pc == done); to next−list
```

```
generateVarNextList(v,S)
 foreach statement s with target(s) = v
  append (pc ==label(s)) ? expression(s) : to
      v−expr
 endfor
 //otherwise v stays the same
 append v; to v−expr
 let v−next−list be v = v−expr ;
```

```
generateDeclarationList(S)
 foreach variable v in V ∪ A
  append typeof(v) v; to decl−list
  // handle variables that are not assigned
  if v is used before it is assigned
   append primary typeof(v) vnondet; to decl−list
  endif
 endfor
 append int pc;
        bool notdone; to decl−list
```

```
generateInitList(S)
 foreach variable v in V ∪ A
  //variables used before assignied are
      nondeterministic
  if v is used before it is assigned
   append v = vnondet; to init−list
  else
   append v = 0; to init−list
  endif
 endfor

 // initialize program counter
 append pc =label(first(S));
        notdone = true; to init−list
```

Fig. 5: Algorithms for generating the $\mathcal{O}lp$ `decl-list`, `init-list`, and `next-list` from $\mathcal{S}$

```
generateArrayNextList(v,S)
 //iterate over all statements assigning to v
 foreach assignment s with base(target(s)) = v
  //s is of the form v[ie] = e
  let ie be index(target(s))
  let e be expr(s)
  //aggregate array access in one expression
  append (pc ==label(s)) ? ie : to a−index
  append (pc ==label(s)) ? e : to a−expr
 endfor
 //when no array access happens, v[0] does not
     change
 append 0 to a−index
 append v[0] to a−expr
 //construct one assignment statement for the
     array
 let v−next−list be v[a−index] =a−expr;
```

```
generatePCNextList(S)
 // encode control flow into pc
 let pc−next−list be pc =
 // iterate over all statements
 foreach s in S
  if s is a conditional //handle conditional
   append (pc ==label(s)) ?
            (condition(s) ? then(s) : else(s)) :
          to pc−next−list
  elseif s is a loop     // handle loops
   append (pc ==label(s)) ?
            (condition(s) ? body(s) : next(s)) :
          to pc−next−list
  elseif s is an assignment with target(s) = pc
   // handle direct assignment to pc
   // which result from function call resolution
   append (pc ==label(s)) ? expr(s): to pc−next−list
  else // handle other statements (assignments)
   append (pc ==label(s)) ? next(s): to pc−next−list
  endif
 endfor
 // when done, program counter does not change
 append pc; to pc−next−list
```

Fig. 6: Algorithm to aggregate array variable assignments (left) and encode control flow into program counter pc (right)

The `generateOneLoopProgram` Algorithm builds one assignment statement `v-next-list` for each variable v and appends it to `next-list`. If v is a regular variable, `generateVarNextList` iterates over each assignment s where v is the target and builds a nested ternary conditional expression `v-expr` that evaluates to `expr(s)` when pc points to the label of s. When pc does not point to a statement that assigns to v, `v-expr` evaluates to the original value of v which is is denoted by appending v as the value of the last choice in in the nested ternary expression.

If v is an array variable, Algorithm `generateArrayNextList` from Figure 6 iterates over all statements s where v is target. It aggregates all expressions and index expressions into two nested ternary conditional expressions `a-index` and `a-expr` for the index expression and the right hand side expression of the `v-next-list` assignment statement. The ternary conditional expressions depend on the position of pc to return the corresponding expressions and index expressions. In case pc does not point to a statement with v as a target, then `v-next-list` makes sure that $v[0]$ stays the same by appending 0 and $v[0]$ as the last default choices in the ternary conditional expressions, respectively.

Finally, `generateOneLoopProgram` calls `generatePCNextList` from Figure 6 to encode the control flow of $\mathcal{S}$ into `pc-next-list` with a nested ternary conditional expression that defines the value of pc. Algorithm `generatePCNextList` iterates over all statements and encodes their execution flow semantics. When s is at a conditional statement, pc moves to `then-block` if the Boolean condition evaluates to $true$, and to the `else-block` otherwise. When s is at a loop statement, then pc moves to the first statement of `body-block` if the Boolean condition evaluates to $true$, and to the statement next to the loop otherwise. When s is an assignment statement such that `target(`s`)`$= pc$, then pc moves to the expression of s; this is similar to a `goto` statement and may result from resolving function calls as discussed in Section 4.3. For other statements, pc points to the next statement.

4.3 Preprocessing J programs

$\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ resolves function declarations, return statements, functions calls, pre and post conditions, and quantifiers by translating $\mathcal{S}$ into $\mathcal{S}'$ that is restricted to $Jcore$.

Non recursive functions: Algorithm `resolveNonRecursiveFunction` in Figure 7 considers the declaration and function calls of a function $fname$. It declares two additional variables in the function. Variable $return\text{-}pc$ is added to `arg-decl-list` and it holds the label that pc should return to after the function is done. Variable $retvar$ is added to `func-decl-list` and it holds the value of the return expression when the function is done. The `func-return-statement` is then replaced by `func-return-list` that sets $retvar$ and pc to `expr(func-return-statement)` and $return\text{-}pc$, respectively.

For each statement s containing an $fname$ function call, `resolveNonRecursiveFunction` constructs a list of assignment statements that marshal the arguments and set the $return\text{-}pc$ to

```
resolveNonRecursiveFunction(fname,S)
 let S be of the form decl-list statement-list
 let type fname(arg-decl-list) {
       func-decl-list func-body func-return-statement }
   be the non-recursive function declaration of fname

 // add a return program counter variable to function argument
 // declarations to hold the label where the function should return
 append , int return-pc to arg-decl-list
 // add a return variable to the function variable declaration
 append type retvar; to func-decl-list
 // turn the return statement into an assignment statement
 append retvar = expression(func-return-statement); to func-return-list
 // when function done, move program counter back to the caller
 append pc = return-pc; to func-return-list
 replace func-return-statement with func-return-list
 // iterate over all statements with calls to fname
 foreach statement s containing an fname(arg-call-list) function call
  // marshal the argument list
  foreach argument expression e in arg-call-list
   let a be the corresponding variable to e in arg-decl-list
   append a = e; to func-call-list
  endfor
  // save the label of the statement the function should return to
  // and transfer control to the function
  append return-pc = label(fcall-ret-statement);
         pc = label(first(func-body)); to func-call-list
  // declare a unique function call return variable
  let i be a unique number
  append type fcall-retvar-i; to decl-list
  // create a statement that reads the return value from the function
  let fcall-ret-statement be fcall-retvar-i = retvar;
  // place the return point from the function
  append fcall-ret-statement to func-call-list
  // fix the statement to use the returned value
  s' = substitute fname(arg-call-list) by fcall-retvar-i in s
  append  s' to func-call-list

  replace s with func-call-list
 endfor
```

Fig. 7: Resolve non-recursive function calls and declarations

the label of `fcall-ret-statement`. It then adds an explicit statement that sets pc to the label of the first statement in $fname$. The `fcall-ret-statement` reads the value of $retvar$ into a unique function call return variable $fcall\text{-}retvar\text{-}i$. This is necessary to resolve multiple function calls in the same expression. Statement s' substitutes the function call by $fcall\text{-}retvar\text{-}i$. The argument marshaling, `fcall-ret-statement`, and s' form the list `func-call-list`. Finally, statement s is replaced by the list of statements `func-call-list`.

Recursive functions: Algorithm `resolveRecursiveFunctionDeclaration` of Figure 8 resolves a recursive function declaration $fname$ into *Jcore* by emulating stack frames with depth $max\text{-}depth$. The argument and local variables of $fname$ become arrays. References to them in the body of the function are replaced with the corresponding array variables indexed by an additional stack pointer $sp\text{-}i$ that points to the current recursive depth of the function. The rest

```
resolveRecursiveFunctionDeclaration(fname,S)
 let S be of the form decl-list statement-list
 let type fname(arg-decl-list) {
       func-decl-list func-body func-return-statement }
   be the recursive function declaration of fname
 //increase the dimensionality of arguments
 foreach argument declaration arg-decl in arg-decl-list
  // type var becomes type var[max-depth] and
  // type array-var[c] becomes type array-var[c][max-depth]
  append arg-decl[max-depth], to arg-decl-list-sp
 endfor
 // add a return-pc variable to the argument list
 append int return-pc[max-depth] to arg-decl-list-sp
 replace arg-decl-list with arg-decl-list-sp
 // similarly, increase dimensionality for local variables
 foreach declaration decl in func-decl-list
  append decl[max-depth]; to decl-list-sp
 endfor
 // add return variables
 append int retvar[max-depth]; to decl-list-sp
 replace func-decl-list with decl-list-sp

 // variable  sp-i plays a stack pointer role for fname
 let i be a unique number
 append int sp-i; to decl-list

 // starts at negative 1
 // once incremented it is at the first recursive frame
 prepend sp-i = -1; to statement-list

 // replace variables with frame variables
 foreach variable v in arg-decl-list and func-decl-list
  replace all occurrences v with v[sp-i] in func-body
  replace all occurrences v with v[sp-i] in func-return-statement
 endfor
 // similarly for array variables
 foreach array variable a in arg-decl-list and func-decl-list
  replace all occurrences a[c1] with a[c1][sp-i] in func-body
  replace all occurrences a[c1] with a[c1][sp-i] in func-return-statement
 endfor
 // fix the return statement to return to the callee
 append retvar[sp-i] = expression(func-return-statement);
        pc = return-pc[sp-i]; to func-return-list
 replace func-return-statement with func-return-list
```

Fig. 8: Resolve recursive function declarations

of the translation follows similarly to Algorithm `resolveNonRecursiveFunction` of Figure 7 and declares additional variables to hold the return, and return program counter values.

Algorithm `resolveRecursiveFunctionCalls` of Figure 9 works similarly to resolving function calls in Figure 7 with additional attention to the stack pointer $sp\text{-}i$. Variable $sp\text{-}i$ is used to appropriately marshal the arguments and the return program counter to the current frame of the recursive function.

Specifications and quantifiers: Algorithm `resolveQuantifier` of Figure 10 takes a quantifier expression q and the statement s that contains it. It declares a Boolean variable $q\text{-}i$ in the scope of the quantifier to hold the value of the quantifier. It initializes $q\text{-}i$ to $true$ when q is a universal

```
resolveRecursiveFunctionCalls(fname, S)
 let S be of the form decl-list statement-list
 let type fname(arg-decl-list) {
       func-decl-list func-body func-return-statement }
   be the recursive function declaration of fname
 let sp be the stack pointer corresponding to fname in decl-list

 foreach statement s containing an fname(arg-call-list) function call
  // declare a unique function call return variable
  let i be a unique number
  append type fcall-retvar-i; to decl-list
  // create a statement that reads the return value from the function
  append fcall-retvar-i = retvar[sp]; to fcall-ret-statement
  // marshal the argument list
  foreach argument expression e in arg-call-list
   let a be the corresponding variable to e in arg-decl-list
   append a[sp + 1] = e; to func-call-list
  endfor
  // save the label of the statement the function should return to
  append return-pc[sp + 1] = label(fcall-ret-statement); to func-call-list
  // increment the stack pointer
  append sp = sp + 1 to func-call-list
  // transfer control to the function
  append pc = label(first(func-body)); to func-call-list
  // place the return point from the function
  append fcall-ret-statement to func-call-list
  // decrement the stack pointer
  append sp = sp - 1 to func-call-list
  // fix the statement to use the returned value
  s' = substitute fname(arg-call-list) by fcall-retvar-i in s
  append  s' to func-call-list

  replace s with func-call-list
 endfor
```

Fig. 9: Resolve recursive function calls

quantifier, and to *false* when q is an existential quantifier. It translates the quantifier statement to a loop that iterates over the range of the quantified variable v and accumulates the value of the Boolean expression with a conjunction in case q was a universal quantifier, and a disjunction in case q was an existential quantifier. Algorithm `resolveQuantifier` substitutes for q in s by q-i to construct s'. Finally, it replaces s with the constructed loop and s'. Alternatively, if the Boolean expression in the quantifier does not include function calls, the quantifier can be unrolled into a sequence of conjunctions (disjunctions in case of an existential quantifier) of the Boolean expression for the range of v. This allows for a smaller execution time.

Algorithm `resolvePrePost` of Figure 10 declares a Boolean variable `pre-specname` for each precondition statement s of the form `@pre specname {boolean-expr}`. It then replaces the precondition statement with an assignment statement that updates `pre-specname` with the value of the corresponding Boolean expression `boolean-expr`. Algorithm `resolvePrePost` works similarly for the postcondition. Of special interest is the expression $\mathit{pre\text{-}specname} \wedge pc = \mathit{label}(s) \rightarrow \mathit{post\text{-}specname}$ which expresses that the program satisfies the specification.

```
resolveQuantifier(q, s, S) // statement s contains quantifier q
 let q be of the form Q (v [e1 ... e2]) (boolean-expr)
 let scope-decl-list and scope-statement-list be the declaration
     and statement lists of the scope of s, respectively
 let i be a unique number // create a unique Boolean variable
 append bool q-i; to scope-decl-list
 append v = e1; to q-list // initialize the quantified variable
 if Q is a universal quantifier
  prepend q-i = true; to scope-statement-list // initialize q-i
  // translate the quantifier into a loop
  append while(v ≤ e2 ∧ q-i){q-i = q-i∧ boolean-expr; v = v + 1;} to q-list
 else // Q is an existential quantifier
  prepend q-i = false; to scope-statement-list
  append while(v ≤ e2 ∧ ¬q-i){q-i = q-i∨ boolean-expr; v = v + 1;} to q-list
 endif
 s' = substitute q by q-i in s // substitute back into the code
 append s' to q-list
 replace s with q-list
```

```
resolvePrePost(S)
 //declare a Boolean variable for each pre/
     post condition and replace the
     statement by an assignment to the
     declared variable
 foreach precondition statement s of the
     form (@pre specname{boolean-expr})
  append bool pre-specname; to decl-list
  replace s with pre-specname =boolean-expr;
 endfor
 foreach postcondition statement s of the
     form (@post specname{boolean-expr})
  append bool post-specname; to decl-list
  replace s with post-specname =boolean-expr;
 endfor
```

Fig. 10: Resolve quantifiers (left) and pre- and post condition statements (right)

```
traverse(exp)
 // base case of recursion
 if (exp is a variable)
  return vargates(exp)
 elseif (exp is constant array access)
  return vargates(exp)
 elseif (exp is array access)
  return resolveArrayAccess(exp)
 endif
 // traverse operands and store
     results in wirevec
 for i from 1 to exp.operands.size()
  wirevec[i] = traverse(exp.operands[i])
 endfor
 //lookup AIG circuit for operation in
     library
 return library(exp.operation, wirevec)
```

```
variables(decl-list)
 foreach variable v in decl-list
  if v is an array variable
   //handle array variables
   for j from 0 to size(v) - 1
    vargates(v[j]) =
      instantiate-registers(v[j], type(v[j]))
   endfor
  elseif (v is a primary input)
   // handle nondeterministic variables
   vargates(v) =
     instantitate-primary-inputs(v, type(v))
  else
   vargates(v) =
     instantiate-registers(v, type(v))
  endif
 endfor
```

Fig. 11: Expression traversal and register instantiations for variables

4.4 $\mathcal{O}lp$ to AIG circuits

Algorithm `variables` of Figure 11 instantiates vectors of AIG registers and primary inputs to translate the corresponding $\mathcal{O}lp$ program variables and stores the translation in a lookup function `vargates`. The width of a bit vector can be selected by the user, or set to match the default width of the declared type. Typically the default values for the bit width are 32 bits for an integer. Arrays are represented by a fixed number of array elements and each element is then treated as a regular variable. The number of array elements is either specified in the

program, bounded by the user, or fixed to a constant by $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$.

4.4.1 Assignment statements: $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ considers each assignment statement s in `init-list` and `next-list` and traverses the right hand side expression of s with the recursive `traverse` Algorithm of Figure 11. If expression exp refers to a variable (base case), or an array access operator with a constant index, then `traverse` returns `vargates(`exp`)`. If exp was an array access $a[ie]$ where a is the array variable and ie is an index expression, then `resolveArrayAccess` of Figure 12 translates exp into a nested ternary expression where the conditions depend on the value of ie and the values are array access terms with constant indices between 0 and $size(a)-1$. Algorithm `resolveArrayAccess` calls `traverse` again on the generated expression to produce the corresponding AIG.

If the expression is a logical, conditional, or arithmetic expression, then the `library` routine looks it up in a complete table of AIG circuits with the adequate bit width. For example, if the expression is a ternary conditional statement of the form $b\ ?\ e_1 : e_2$, then `library` instantiates a multiplexer, connects its two data fanins to the nodes corresponding to e_1 and e_2, connects its control fanins to the nodes corresponding to b, and returns its fanouts. Alternatively, users have the choice to abstract multiplication, division, and remainder by non-constant factors with non-deterministic variables (uninterpreted functions).

4.4.2 Connections and array access target terms: For assignment statements with target terms that are regular variables, or array access with constant index expressions, $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ connects the nodes corresponding to the right hand side expressions in `init-list` and `next-list` to the initial and next state value fanins of the corresponding register gates, respectively.

An assignment statement s of the form $a[ie] = expr;$, where a is an array variable, ie is a non constant index expression, and $expr$ is an expression, requires preprocessing before being translated to AIG. Algorithm `variables` instantiates $size(a)-1$ register vectors corresponding to the elements of array variable a, each requiring initial and next state value functions. Algorithm `resolveArrayTargetTerms` of Figure 12 takes as input an assignment statement s and replaces it with $size(a)-1$ assignment statements of the form $a[j] = (j == ie)?expr : a[j];$ where j ranges from 0 to $size(a)-1$ such that each array element $a[j]$ evaluates to $expr$ if ie evaluates to j, otherwise, $a[j]$ keeps its value.

```
resolveArrayAccess(exp)
 let a be base(exp)
 let ie be index(exp)
 let N be size(a) - 1

 for j from 0 to N − 1
  append (ie == j) ? a[j] : to
      a−expr−resolve
 endfor
 append a[N] to a−expr−resolve
 return traverse(exp)
```

```
resolveArrayTargetTerms(S)
 foreach statement s where target(s) is array−access
  let a be base(array−access)
  let ie be index(array−access)
  let N be size (a) - 1

  for j from 0 to N
   append a[j] = (j == ie) ? expr : a[j]; to a−resolve
  endfor
  replace s with a−resolve
 endfor
```

Fig. 12: Resolve array access expressions and resolve array target terms

5 Implementation

We fully implemented $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ using C++ and ANTLR [41] and integrated our implementation with the ABC synthesis and verification framework [7] and the GTKWave visualization tool [8]. The frontend supports the J syntax under C, C++, and Java. Figure 13 shows a trace that highlights the defect of the array search program of Figure 4 where d is not in a and rv is not -1. $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ provides a fully automated verification path of execution where it executes a well selected sequence of synthesis reduction and verification techniques that work well for software verification tasks. $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ also supports an interactive shell with parser, synthesis, verification, and debugging commands where the user can save and use intermediary results, try different synthesis and verification strategies, and inspect counterexamples. Table 2 provides a short list of the $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ commands. The commands allow the users to specify (1) the program, (2) bounds on variable data width, array size, recursion depth, and (3) optional automatically embedded checks on array out of bound access and arithmetic overflow. $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ is available online as an open source tool with tutorial and documentation [2].

6 Results

We evaluated $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ by verifying software artifacts including library algorithms and data structures with complete sophisticated specifications from the Calculus of Computation book [42], an array based implementation of the red black balanced binary search tree (RBBST)with specifications formalized from [43], and a memory allocation model (MMAN) with complete specifications provided from the Frama-C ANSI C Specification Language (ACSL) [44]. Table 3 reports

2. http://research-fadi.aub.edu.lb/dkwk/doku.php?id=sa

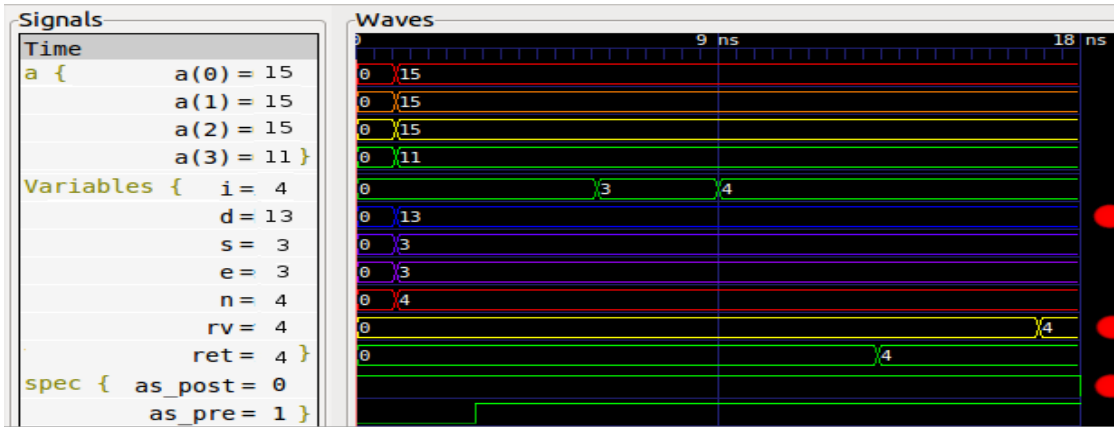

Fig. 13: Array search visual debugging with GTKWave [8].

TABLE 2: Summary of $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ commands

Command	Description
`read`	Parses a program and generates a parse graph.
`print-time`	Prints the runtime time of the last command and the total run time.
`prove-scheme1`	Performs a predefined scheme of sequential synthesis and proof strategies.
`prove-scheme2`	Performs a predefined scheme of sequential synthesis and proof strategies.
`abc cmd`	Performs the ABC techniques specified in `cmd`.
`debug`	Generates a program counterexample from the AIG counterexample and allows the user to inspect the values in a textual format.
`vdebug`	A visual version of `debug` that allows the user to view and interact with the traces using GTKWave.
`start-sim`	Starts the simulator helper tool to help the user analyze the code and perform what-if analysis on variable values.
`sim-variable`	Runs the simulator helper tool to generate and display values for specific variables.
`start-prove`	Configures the prove engine with parameters such as variable bit width, array size bounds, and recursion depth bounds, as well as overflow check, out of bound access check, and quantifier unrolling.

the results with several bounds and compares against CBMC [6]. The results show that $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ scales to bounds larger than CBMC by at least one order of magnitude.

We also evaluated $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ with benchmark verification tasks from the 2014 software verification competition [21] each with memory safety, termination, and error label reachability checks. Some of the tasks are *true* where a proof is expected, and some are *false* where a counterexample is expected. We ran our experiments on a machine with similar settings to the server machine reported in [21] (3.4 GHz 64-bit Quad Core CPU, a 64 bit GNU/Linux operating system, and 32 GB of RAM). We compare our results to the leading tools in the competition. Table 4 reports the results where some of the leading tools in the competition (CBMC [6], ESBMC [45], BLAST [46], and CPAChecker [47]) successfully completed the verification task and compares against them in terms of running time. The results show that $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ succeeded to verify all the tasks and ranked either first or second in terms of running time in comparison

with the leading tools. Table 5 reports the results of verifying benchmarks from [21] where the leading tools reported a timeout and failed to complete the verification task. According to the wrapper scripts provided from [21], CBMC and ESBMC report a $true$ result by default when they reach their timeout while the SAT or SMT solver are running, they report an $unknown$ result when the timeout passes before they generate the first CNF or SMT instance to pass to the underlying solver. This is *probably reasonable* as CBMC and ESBMC consider the inability of the solvers to find counterexamples within the timeout period as an indication of the rarity or absence of such results. Thus we consider a $true$ result with a timeout time from CBMC and ESBMC as a timeout result for comparison purposes. The results show that $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ succeeded to complete verification tasks that were not completed by the leading tools.

In the experiments, we used an automated script that ran reduction algorithms first with a timeout of 10 minutes with Table 4 and three minutes for Tables 4 and 5, followed by three instances of verification using the `pdr`, `dprove`, and `ind` ABC commands with proper timeout, number of frames, and BDD size configurations.

6.1 Algorithms and data structures

Table3 shows that $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ was able to verify programs with sophisticated specifications faster than CBMC and for array sizes that CBMC could not verify. The table shows the number of memory elements (registers), gates, and logic levels before and after the ABC reductions. It also shows the total time taken and the time taken to verify the programs. The CBMC columns show the size of the generated CNF formulae in terms of the number of CNF variables and clauses. We developed the programs and the specifications and called $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ and CBMC to verify their correctness incrementally. Both $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ and CBMC returned counterexamples that we used to correct the programs and the specifications. For programs with array indirection where array elements are used as indexes of other arrays such as `next`, and `left`, `right`, and `parent` in the linked list and red black balanced binary search tree (RBBBST), $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ and CBMC both succeeded in returning counterexamples for index out of bound violations. However, once we corrected those issues, CBMC mistakenly reported that the programs are correct while $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ correctly reported counterexamples related to the cardinality of the data structures after insertion or removal of elements. The rows highlighted in red report the time taken by CBMC to verify the defected code and return a wrong result versus the time taken by $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ to verify the correct code. Note that CBMC and other tools reportedly returned several wrong results for

TABLE 3: Results of $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ and CBMC to verify array based algorithms and data structures with increasing array sizes. TO, MO, ERR, MISTAKE stand for timeout, memory out, runtime error, and verification mistake, respectively.

	Array	AIG size before/after reductions			$\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$		CNF size		CBMC
	size	registers	and	level	verif.	total (s)	vars	clauses	time (s)
array search	3	86 / 41	719 / 313	24 / 15	0.33	4.36	2,416	6,784	0.016
	7	118 / 68	1,064 / 568	27 / 19	3.89	12.4	4,612	15,008	722.4
	15	174 / 119	1,781 / 1,116	30 / 21	2.41	16.87	9,112	34,496	TO
	31	286 / 226	3,362 / 2,346	45 / 24	1.43	33.67	18,332	84,928	TO
	63	526 / 461	6,895 / 5,100	78 / 26	4.57	99.64	37,216	230,208	TO
	127	1,054 / 984	14,780 / 11,315	143 / 28	21.32	397	75,876	695,616	TO
	255	4,798 / 4,718	70,742 / 55,364	529 / 33	682.1	8,022.1	TO	TO	TO
binary search	3	94 / 56	879 / 555	30 / 19	0.11	1.04	6,503	24,533	0.085
	7	151 / 83	1,832 / 850	42 / 17	0.54	1.47	16,172	68,130	1.91
	15	268 / 143	5,185 / 1,943	62 / 20	25.42	27.69	42,461	197,223	38.493
	31	491/262	18,355/4,903	68/28	112.84	115.13	100,379	412,410	4,955.223
	63	874/454	64,457/8,608	102/28	1,132	1,152.22	2,660,677	13,689,632	TO
	127	1,911/976	277,662/19,907	170/32	TO	TO	TO	TO	TO
bubble sort	3	114 / 44	1,198 / 393	29 / 16	0.29	5.79	15,534	43,332	0.174
	7	169 / 68	2,218 / 885	35 / 20	17.1	31.09	63,785	197,603	7.298
	15	276 / 117	5,607 / 2,106	47 / 22	1,390.25	1,426.98	341,552	1,082,480	455.5
	31	603/360	40,448/25,546	75/36	2,668.22	2,669.82	2,144,801	6,809,622	TO
	63	1,293/814	174,535/110,123	140/40	4,323.51	4,384.08	TO	TO	TO
	127	2,847/1,844	770,398/485,523	269/44	TO	TO	T0	TO	TO
selection sort	3	86 / 45	1,021 /679	26 / 19	20.62	20.69	17,941	55,151	0.37
	7	150 / 87	2,726 / 1,971	34 / 24	1,105.2	1,105.74	315,505	1,129,784	TO
	15	280 / 125	5,676 / 2,236	47 / 22	2,280	2,280.18	959,841	3,849,448	TO
	31	590/383	39,543/26,547	75/27	TO	TO	3,269,941	14,305,256	TO
	63	1,278/841	172,385/112,601	140/31	TO	TO	TO	TO	TO
array partition	3	110 / 57	1,896 / 689	33 / 19	0.87	2.79	36,636	121,478	2.43
	7	147 / 87	2,509 / 1,174	34 / 24	93.47	97.56	622,799	2,663,342	TO
	15	206/ 138	3,819 / 2,651	38 / 35	2,127	2,135.7	2,061,314	9,390,063	TO
	31	323/248	6,779/5,557	42/39	TO	TO	7,315,089	34,755,647	TO
	63	568/486	13,451/12,178	46/43	TO	TO	TO	TO	TO
linked list insert	3	237 / 98	4,310 / 871	38 / 19	109.63	118.89	28,347	28,347	0.194 (mistake)
	7	344 / 179	6,117 / 1,693	41 / 26	1,800	1,811	83,079	83,079	0.527 (mistake)
	15	503/345	10,521/3,763	47/26	TO	TO	281,759	281,759	2.027 (mistake)
	31	839/634	19,624/10,060	53/36	TO	TO	1,039,655	1,039,655	13.575 (mistake)
	63	1,559/1,482	40,375/16,396	59/32	error	error	3,946,601	3,946,601	152.92 (mistake)
linked list remove	3	197 / 84	2,906 / 722	33 / 21	47.72	71.383	22,898	22,898	0.163 (mistake)
	7	293 / 157	4,454 / 1,387	39 / 25	1,800.15	1,830.21	59,219	59,219	0.457 (mistake)
	15	448/397	7,847/3,094	45/20	TO	TO	175,176	175,176	1.935 (mistake)
	31	778/821	15,763/6,364	51/23	TO	TO	579,957	579,957	11.163 (mistake)
	63	1,492/1,584	233,791/14,346	57/25	TO	TO	2,038,358	2,038,358	85.94 (mistake)
RBBBST insert	3	779/407	14,284/6,451	122/41	313.11	314.22	2,765,445	2,765,445	7.169 (mistake)
	7	1,555/1,077	40,794/29,281	128/54	3,600.48	4,087.1	5,602,125	5,602,125	15.456 (mistake)
	15	1,834/1,221	50,540/28,928	134/43	TO	TO	10,487,436	38,445,954	TO
	31	2,810/2,124	100,841/66,558	152/49	TO	TO	21,680,107	21,680,107	TO
	63	4,789/4,034	213,576/147,577	186/56	TO	TO	MEMOUT	MEMOUT	MEMOUT
RBBBST remove	3	1,487/923	34,506/15,309	251/44	124.63	313.51	4,403,782	4,403,782	8.57 (mistake)
	7	1,559/993	43,608/18,963	253/41	3,600.62	4,154.93	8,587,721	8,587,721	17.972 (mistake)
	15	2,072/1,404	69,646/48,341	254/56	TO	TO	17,476,396	17,476,396	44.67 (mistake)
	31	2,969/2,308	139,461/101,039	259/53	TO	TO	37,383,755	37,383,755	327.682 (mistake)
	63	4,610/3,901	291,311/233,674	262/59	TO	TO	MEMOUT	MEMOUT	MEMOUT
memory manager	16	3,033/962	41,137/13,461	366/31	460.09	608.61	TO	TO	TO
	64	8,799/2,873	99,510/52,543	1,072/36	1,759.52	2,238.18	TO	TO	TO
average reductions		64.5%	51.2%	44.7%					

TABLE 4: Runtime results in seconds for verification tasks completed successfully by $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ compared to the leading software verification competition tools. TO and ERR stand for timeout and error, respectively.

	Category	$\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$	BLAST	CBMC	CPAChecker	ESBMC
Problem01_00_true	eca	7.5	TO	TO-T	4.5	4.6
Problem01_10_true	eca	8.59	730	TO-T	4.5	4.2
Problem01_30_true	eca	9.06	TO	TO-T	4.6	4
Problem01_40_true	eca	9.41	TO	TO-T	4.5	5.9
Problem02_00_true	eca	10.68	460	TO-T	4.3	6.3
Problem02_10_true	eca	9.71	580	TO-T	4.3	4.8
Problem02_20_true	eca	10.14	ERR	TO-T	4.3	4.8
Problem01_20_false	eca	9.35	TO	150	5.3	33
Problem01_50_false	eca	9.87	TO	22	4.8	33
Problem01_60_false	eca	5.38	780	0.56	2.8	30
s3_clnt_1_true	ssh-simplified	57.3	140	88	3.5	3.4
s3_clnt_1_false	ssh-simplified	47.15	23	2.1	2.8	13
s3_clnt_2_true	ssh-simplified	47.1	TO	90	3.6	4.2
s3_clnt_2_false	ssh-simplified	44.26	23	2.1	2.8	14
s3_clnt_3_true	ssh-simplified	77.64	350	110	3.7	3.9
s3_clnt_3_false	ssh-simplified	27.56	23	2.5	2.8	17
s3_clnt_4_true	ssh-simplified	79.45	TO	90	3.5	3.6
s3_clnt_4_false	ssh-simplified	76.95	23	2.1	2.8	14
s3_srvr_1a_true	ssh-simplified	20.11	1.2	3.4	1.8	0.69
s3_srvr_1b_true	ssh-simplified	1.68	1.2	1.3	1.6	0.48
s3_srvr_1_false	ssh-simplified	35.47	4.2	3	2.4	22
s3_srvr_2_true	ssh-simplified	567.51	410	200	4.4	3.9
s3_srvr_2_false	ssh-simplified	37.05	4.4	2.8	2.4	18
s3_srvr_4_true	ssh-simplified	799	440	210	14	3.8
s3_srvr_6_true	ssh-simplified	694	TO	230	24	4.3
s3_srvr_6_false	ssh-simplified	14.24	0.15	0.21	19	21
s3_srvr_8_true	ssh-simplified	340	220	220	4.7	4.1
test_locks_5_true	ControlFlowInt	0.15	23	1.6	2.1	0.43
test_locks_6_true	ControlFlowInt	0.23	44	2.4	2.7	0.42
test_locks_7_true	ControlFlowInt	0.31	140	2.8	5	0.45
test_locks_8_true	ControlFlowInt	0.37	280	3.6	21	0.48
test_locks_9_true	ControlFlowInt	0.55	610	4.5	53	0.52
test_locks_10_true	ControlFlowInt	0.84	TO	5.6	53	0.54
test_locks_11_true	ControlFlowInt	1.24	TO	7	53	0.6
test_locks_12_true	ControlFlowInt	0.99	TO	8.7	53	0.62
test_locks_13_true	ControlFlowInt	1.14	TO	11	52	0.66
test_locks_14_true	ControlFlowInt	1.56	TO	14	52	0.68
test_locks_15_true	ControlFlowInt	1.83	TO	15	53	0.77
test_locks_14_false	ControlFlowInt	0.46	0.14	0.22	2	1.4
test_locks_15_false	ControlFlowInt	0.51	0.13	0.21	2	1.5

benchmarks in the 2014 software verification competition [21], $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ performed better than CBMC as follows.

- Most of the times where both $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ and CBMC succeeded to verify the programs, $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ took less time to complete the verification task.
- $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ succeeded to verify the programs for bounds larger than what CBMC could verify

TABLE 5: Results for verification tasks in seconds where $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ terminated successfuly while the leading tools either timed out or produced erroneous output. TO: timeout, TO-T: report true on timeout, TO-F report false on timeout, MISTAKE: report erroneous results, ERR: reported a runtime error, and UNKNOWN: report inconclusive result.

	Category	$\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$	BLAST	CBMC	CPAChecker	ESBMC
Problem07_00_true	ControlFlowInt/eca	620	TO	TO-T	TO	UNKNOWN-TO
Problem09_00_true	ControlFlowInt/eca	433	TO	UNKNOWN-TO	TO	UNKNOWN-190
Ackermann01_true	Recursive	723	N/A	TO-T	ERR-1.9	MISTAKE-TO
Ackermann02_false	Recursive	22	N/A	0.73	ERR-1.9	TO-F
Addition03_false	Recursive	11.2	N/A	MISTAKE-TO	ERR-2	TO_F
EvenOdd01_true	Recursive	0.68	N/A	0.96	ERR-1.9	MISTAKE-0.4
EvenOdd03_false	Recursive	0.45	N/A	0.2	ERR-1.9	0.4
insertion_sort_true	loops	154	MISS-6.6	TO-T	TO	UNKNOWN-TO
invert_string_false	loops	0.22	1.8	0.47	TO	0.6
jain_5_true	bitvector	1.44	N/A	0.53	TO	0.25
linear_search_false	loops	0.39	0.13	0.62	ERR-2.6	MISTAKE-0.3
McCarthy91_false	Recursive	1.2	N/A	0.19	ERR-1.9	TO-F
McCarthy91_true	Recursive	567	N/A	TO-T	ERR-1.9	MISTAKE-TO
MultCommutative_true	Recursive	780	N/A	TO-T	ERR-1.9	MISTAKE-TO
s3_srvr_1_alt_true.BV	bitvector	TO	N/A	TO-T	TO	TO-T
s3_srvr_1_true	ssh-simplified	TO	UNKNOWN-0.1	160	4.6	3.9
s3_srvr_3_true	ssh-simplified	TO	430	210	13	4
s3_srvr_7_true	ssh-simplified	TO	TO	210	24	4.3
Problem08_00_true	ControlFlowInt/eca	TO	TO	UNKNOWN-TO	TO	UNKNOWN-TO

within the three hours seconds timeout.

- CBMC failed to generate the CNF formula for some programs with relatively larger array sizes while $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ succeeded to verify some of them and generated AIG circuits but timed out while verifying the rest. This is significant since model checking *typically* returns counterexamples fast when they exist, and takes long time to report a proof for unsatisfiable (correct) claims. That is, $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ succeeded to perform model checking for a significant time length, even when it timed out, as compared to CBMC that failed to generate the CNF formula, thus did not really test the program for correctness.

For the array search program, the precondition checks that the indexes are in range and the postcondition checks whether the return value is consistent. For the binary search program both the precondition and the postcondition specify the sortedness property for the array. For the bubble and selection sort programs, the precondition checks the sanity of the size of the array, and the postcondition specifies that the resulting array holds the same elements of the original array in order. For the array partition program, the precondition specifies the sanity of the indexes and the size of the array, and the postcondition specifies that the resulting array holds the same elements of the original array partitioned in order left and right of the cursor. The linked list insert and remove precondition specifies a non circular, in order, list of elements,

and the postconditions reassert the same specification with an additional cardinality condition. The RBBBST insert and remove precondition specifies the full characteristics of an RBBBST and that the index arrays involved all hold valid indexes. The postconditions reassert the same specification with an additional cardinality condition. The RBBBST insert also checks the sanity of the newly allocated node in the tree. The memory manager unit has several multidimensional arrays to bookkeep the use of memory chucks and the precondition and the postcondition assert the sanity of the arrays, and the consistency of the module. We made use of the assume statement `__CPROVER_assume` and loops to implement the preconditions and the quantifiers in CBMC.

The reduction algorithms used by $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ were able to reduce the original problems on average to 64.5% of their memory elements, and to 51.2%, and 44.7% of their logic gate and level counts in often insignificant time. Without these reductions $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ timed out on several benchmarks. These reductions have no counterparts in CBMC.

6.2 SV-COMP 2014 benchmarks

Tables 4 and 5 show the results of $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ compared against the results reported in [21] for the leading participating software verification tools. We selected the benchmarks that required no or minor syntax modifications for the $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ front end and compared the execution time of $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ with the leading tools. The tables show the name of the benchmark, its category and the running time needed by the tools to complete verification with a timeout of 900 seconds as required by SVComp 2014 regulations.

$\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ succeeded to complete the verification tasks in Table 4 and ranked first or second in terms of running time. $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ succeeded to complete the verification tasks in Table 5 where the leading tools either timed out, reported runtime errors, or reported a wrong result (produced a counterexample when the benchmark is correct, or produced a claim of correctness when the benchmark had a known defect).

For the last five verification tasks in Table 5 $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ timed out without producing a conclusive result while some of the leading tools returned the correct result, reported a timeout, or reported an inconclusive result.

6.2.1 Interactive sessions: We inspected the results of the benchmarks where the automated $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ script did not return a conclusive result. We were able to verify several of them

interactively via trying ABC synthesis reduction and verification commands and techniques with different configurations and settings. This is not possible with the tools we compared to.

7 RELATED WORK

ESBMC [45] uses SMT solvers to verify multi-threaded C programs by forcing bounds on the number of context switches, loop iterations, and recursive calls. It also uses an aligned memory model that resolves pointers. We differ in that we use AIG solvers, where we do not need to have bounds for loop iterations, and we plan in the future to extend our approach to concurrent programs by encoding each execution thread as a separate program counter and limiting the number of program counters.

VCC [28] is an industrial strength verification framework for for concurrent low level C programs. It has a ghost type state system that tracks the validity of memory objects; i.e., references to memory objects do not overlap type states. It generates verification conditions that can be checked by Boogie; i.e. a high order logic analyzer. Boogie in turn uses the Z3 SMT solver [48] for automatic verification and Isabelle [49] for interactive verification. FranklinBIT [50] takes a bit-vector program with a verification condition and computes an unsound approximation (not over and not under) of the verification condition. Then it uses a logic solvers (SMT or propositional) to decide the original verification condition strengthened with the inductive unsound approximation.

LLBMC [51] is a bit precise bounded model checker for low level C programs that works on an intermediate assembly representation of the program using an SMT solver. Ultimate Automizer [52] computes *appropriate* abstractions of programs based on statements as alphabet atoms in an automata framework. The work in [53] is implemented within the CPACheker [47] platform to verify *event condition action* (ECA) systems using BDDs. Our method allows the use of symbolic model checking through ABC that makes use of BDD without the need to restrict ourselves to ECA systems. The CPAChecker [47] framework passes a program $\mathcal{S}$ and an invariant Ψ to model checkers for verification. In case a model checker returns an inconclusive result, CPAChecker uses partial results to formulate a state predicate formula Φ where the invariant holds. This incrementally reduces the work of the next model checker to check $\neg\Phi$. In the future, we plan to extend the same idea to the ABC model checker. BLAST [46] is the predecessor of CPAChecker and it uses abstraction and refinement techniques to model check temporal properties of C programs. Similar to CPAChecker is UFO [54] that provides a

framework for exploring abstraction and interpolation techniques for software verification and uses SMT solvers as a backend.

CBMC [55] is a bounded model checker for ANSI-C programs that checks for properties such as pointer safety, array bounds and user assert statements. Given an ANSI-C program and a bound on the range of variables, CBMC produces a Boolean formula in CNF and checks the formula with SAT solvers. CBMC relies mostly on the power and speed of SAT solvers. SAT solvers often face an exponential blow up in the number of possible assignments to the atomic propositions. This problem, known as the *state explosion*, and the large number of variables and clauses used in the CNF encoding, limit the CBMC analysis to relatively small bounds.

The work in [26] uses an encoding similar to that of CBMC, but instead of producing a CNF formula it produces an SMT formula. The SMT formula allows for variables with no range bounds. The loops are still unrolled up to a finite bound and loop completion assertions fire in case the bound imposed on the number of loop iterations was small and the loop guard was still satisfiable. The higher level solver may now decide to reproduce a new SMT formula with bigger loop unrolling bounds and call the SMT solver on the new formula. We differ in that our encoding to AIG circuits requires no bounds on loop iterations. This allows for more succinct representation of programs than with SMT.

Java PathFinder [29] (JPF), is a popular explicit-state model checker for Java programs. It implements a customized Java virtual machine (JVM) that supports state backtracking and allows programs to check properties of a wide range of Java programs without the need to respecify the programs in specialized languages. JPF does not apply any program transformations. JPF also does not scale well in the presence of loops and branches with long running time. JPF operates at the word level, however, it does not make use of program specific properties and exhaustively searches the tree of possible executions. We differ in that we operate at the bit level and we apply reduction techniques that allow backtracking algorithms to run faster as they operate on a smaller state space.

The work of Xie et al [32] translates software artifacts into equivalent semantics that are model-checkable. It applies compositional rules to the translated system to build its formal semantics in the context of message passing systems. We differ in that we use a specific program counter based semantics with a finitization rule to translate the software artifacts to an AIG, a model checkable formalization, where reductions, including compositional ones, can be used.

8 Conclusion

We presented $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ that takes a program and a pair of specifications, translates it into an AIG circuit, reduces the circuit using the ABC synthesis reduction techniques, and then check the circuit for correctness using the ABC verification techniques. $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ scales to bounds larger than that possible with existing tools, and solves verification tasks from the software verification competition benchmarks that leading tools at the competition did not solve within the specified timeout. In the future, we plan to extend $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ to verify concurrent programs using several program counters to represent multi-threaded executions. We also plan to integrate $\{\mathcal{P}\}\mathcal{S}\{\mathcal{Q}\}$ with existing frameworks that benefit from word level optimizations and techniques such as CPAChecker to leverage both word level and bit level techniques as well as cover an extended syntax set through existing front ends.

References

[1] F. Zaraket, A. Aziz, and S. Khurshid, "Sequential circuits for relational analysis," in *International Conference on Software Engineering*, May 2007.

[2] H. Mony et al., "Scalable automated verification via expert-system guided transformations," in *Formal Methods in Computer-Aided Design*, Nov. 04.

[3] E. Torlak and D. Jackson, "Kodkod: A relational model finder," in *Tools and Algorithms for Construction and Analysis of Systems*, March 2007.

[4] N. Sorensson and N. Een, "Minisat v1. 13-a sat solver with conflict-clause minimization," *SAT*, 2005.

[5] F. A. Zaraket, A. Aziz, and S. Khurshid, "Sequential circuits for program analysis," in *ASE*, Nov 2007.

[6] E. Clarke, D. Kroening, and F. Lerda, "A tool for checking ansi-c programs," *Tools and Algorithms for the Construction and Analysis of Systems*, pp. 168–176, 2004.

[7] R. Brayton and A. Mishchenko, "ABC: An academic industrial-strength verification tool," in *Computer Aided Verification*, Springer, 2010.

[8] T. Bybell, "GtkWave electronic waveform viewer," 2010.

[9] A. Kuehlmann and J. Baumgartner, "Transformation-based verification using generalized retiming," in *Computer-Aided Verification*, July 2001.

[10] H. Mony, J. Baumgartner, V. Paruthi, and R. Kanzelman, "Exploiting suspected redundancy without proving it," in *Design Automation Conference*, ACM Press, 2005.

[11] A. Kuehlmann, M. Ganai, and V. Paruthi, "Circuit-based Boolean reasoning," in *Design Automation Conference*, pp. 232–237, June 2001.

[12] P. Bjesse and K. Claessen, "SAT-based verification without state space traversal," in *Formal Methods in Computer-Aided Design*, November 2000.

[13] A. Aziz, T. Shiple, V. Singhal, R. Brayton, and A. Sangiovanni-Vincentelli, "Formula Dependent Equivalence for Compositional CTL Model Checking," *Journal of Formal Methods in System Design*, vol. 21, no. 2, 2002.

[14] P. Bjesse and A. Boralv, "DAG-aware circuit compression for formal verification," in *Int'l Conference on Computer-Aided Design*, Nov. 2004.

[15] K. L. McMillan, *Interpolation and SAT-Based Model Checking*. Springer Berlin / Heidelberg, 2003.

[16] D. Wang, *SAT based Abstraction Refinement for Hardware Verification*. PhD thesis, CMU, May 2003.

[17] I.-H. Moon, G. D. Hachtel, and F. Somenzi, "Border-block triangular form and conjunction schedule in image computation," in *Formal Methods in Computer-Aided Design*, Nov. 2000.

[18] M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik, "Chaff: Engineering an efficient SAT solver," in *ACM Design Automation Conference*, June 2001.

[19] P.-H. Ho, T. Shiple, K. Harer, J. Kukula, R. Damiano, V. Bertacco, J. Taylor, and J. Long, "Smart simulation using collaborative formal and simulation engines," in *Int'l Conference on Computer-Aided Design*, Nov. 2000.

[20] J. Baumgartner, A. Kuehlmann, and J. Abraham, "Property checking via structural analysis," in *Computer-Aided Verification*, July 2002.

[21] D. Beyer, "Status report on software verification," in *Tools and Algorithms for the Construction and Analysis of Systems*, Lecture Notes in Computer Science, 2014.

[22] F. A. Aloul, A. Ramani, I. L. Markov, and K. A. Sakallah, "Solving difficult SAT instances in the presence of symmetry," in *Design automation conference*, ACM Press, 2002.

[23] Z. Fu, Y. Yu, and S. Malik, "Considering circuit observability don't cares in CNF satisfiability," in *Proceedings of Design, Automation and Test in Europe*, pp. 1108–1113, March 2005.

[24] Q. Zhu, N. Kitchen, A. Kuelhman, and A. Sangiovanni-Vincentelli, "SAT sweeping with local observability dontcares," in *Design Automation Conference*, July 2006.

[25] S. A. Edwards, "The challenges of hardware synthesis from C-like languages," in *Design Automation and Test in Europe*, 2005.

[26] R. Arm, J. Mantovani, and L. Platania, "Bounded model checking of software using SMT solvers instead of SAT solvers," in *SPIN. Volume 3925 of LNCS*, pp. 146–162, Springer, 2006.

[27] E. Clarke, D. Kroening, N. Sharygina, and K. Yorav, "SATABS: SAT-based predicate abstraction for ANSI-C," in *Tools and Algorithms for the Construction and Analysis of Systems*, Springer Verlag, 2005.

[28] E. Cohen, M. Dahlweid, M. A. Hillebrand, D. Leinenbach, M. Moskal, T. Santen, W. Schulte, and S. Tobies, "VCC: A practical system for verifying concurrent C," in *Theorem Proving in Higher Order Logics*, Springer, 2009.

[29] W. Visser, K. Havelund, G. Brat, and S.-J. Park, "Model checking programs," in *Automated Software Engineering Journal*, April 2003.

[30] H. Zhang and J. Zhang, *Generating models by SEM*. Springer-Verlag, 1996.

[31] P. Baumgartner, A. Fuchs, H. de Nivelle, and C. Tinelli, "Computing finite models by reduction to function-free clause logic," in *Journal of Applied Logic*, 2007.

[32] F. Xie, J. C. Browne, and R. P. Kurshan, "Translation-based compositional reasoning for software systems," in *Formal Methods Europe*, 2003.

[33] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang, "Symbolic model checking: 10^{20} states and beyond," *Information and Computation*, vol. 98, no. 2, 1992.

[34] A. Mishchenko, M. Case, R. Brayton, and S. Jang, "Scalable and scalably-verifiable sequential synthesis," in *Computer-Aided Design, 2008. ICCAD 2008. IEEE/ACM International Conference on*, pp. 234–241, IEEE, 2008.

[35] P. Bjesse and A. Boralv, "Dag-aware circuit compression for formal verification," in *Proceedings of the 2004 IEEE/ACM International conference on Computer-aided design*, pp. 42–49, IEEE Computer Society, 2004.

[36] A. Mishchenko, S. Chatterjee, and R. Brayton, "Dag-aware aig rewriting a fresh look at combinational logic synthesis," in *Proceedings of the 43rd annual Design Automation Conference*, pp. 532–535, ACM, 2006.

[37] A. P. Hurst, A. Mishchenko, and R. K. Brayton, "Fast minimum-register retiming via binary maximum-flow," in *Formal Methods in Computer Aided Design, 2007. FMCAD'07*, pp. 181–187, IEEE, 2007.

[38] N. Eén and N. Sörensson, "Temporal induction by incremental SAT solving," *Electronic Notes in Theoretical Computer Science*, vol. 89, no. 4, 2003.

[39] N. Amla, X. Du, A. Kuehlmann, R. P. Kurshan, and K. L. McMillan, "An analysis of SAT-based model checking techniques in an industrial environment," in *Correct hardware design and verification methods*, Springer, 2005.

[40] N. Een, A. Mishchenko, and R. Brayton, "Efficient implementation of property directed reachability," in *Formal Methods in Computer-Aided Design (FMCAD), 2011*, pp. 125–134, IEEE, 2011.

[41] T. Parr and R. Quong, "ANTLR: A predicated-LL (k) parser generator," *Software: Practice and Experience*, 1995.

[42] A. R. Bradley and Z. Manna, *The calculus of computation*. Springer, 2007.

[43] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*. MIT Press, Cambridge, MA, 1990.

[44] P. Baudin, J. C. Filliâtre, T. Hubert, C. Marché, B. Monate, Y. Moy, and V. Prevosto, *ACSL: ANSI C Specification Language (preliminary design V1.9)*, May 2009.

[45] J. Morse, M. Ramalho, L. Cordeiro, D. Nicole, and B. Fischer, "ESBMC 1.22 - (competition contribution)," in *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 2014.

[46] D. Beyer, T. A. Henzinger, R. Jhala, and R. Majumdar, "The software model checker Blast: Applications to software engineering," *International Journal of Software Tools and Technology Transfer*, vol. 9, October 2007.

[47] D. Beyer, T. A. Henzinger, M. E. Keremoglu, and P. Wendler, "Conditional model checking: A technique to pass information between verifiers," in *Foundations of Software Engineering*, FSE, 2012.

[48] N. Bjrner and L. de Moura, "Z3 10: Applications, enablers, challenges and directions," in *CFV*, 2009.

[49] T. Nipkow, L. C. Paulson, and M. Wenzel, *Isabelle/HOL — A Proof Assistant for Higher-Order Logic*, vol. 2283 of *LNCS*. Springer, 2002.

[50] A. Gurfinkel, A. Belov, and J. Marques-Silva, "Synthesizing safe bit-precise invariants," in *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pp. 93–108, 2014.

[51] S. Falke, F. Merz, and C. Sinz, "The bounded model checker LLBMC," in *Automated Software Engineering (ASE)*, pp. 706–709, November 2013.

[52] M. Heizmann, J. Hoenicke, and A. Podelski, "Software model checking for people who love automata," in *Computer Aided Verification*, vol. 8044, pp. 36–52, Springer Berlin Heidelberg, 2013.

[53] D. Beyer and A. Stahlbauer, "BDD-Based software model checking with CPAchecker," in *Mathematical and Engineering Methods in Computer Science (MEMICS)*, 2012.

[54] A. Albarghouthi, Y. Li, A. Gurfinkel, and M. Chechik, "UFO: A framework for abstraction and interpolation-based software verification," in *Computer Aided Verification*, 2012.

[55] E. Clarke, D. Kroening, and F. Lerda, "A tool for checking ANSI-C programs," in *Tools and Algorithms for the Construction and Analysis of Systems*, March 2004.